Why is AI for robots so hard?
Where is robotics' ChatGPT moment?

Kai Olav Ellefsen

Group for Robotics and Interactive Systems (ROBIN)

University of Oslo

Image: Eve from 1X robotics

What makes learning robot skills harder than learning to master language (and vision)?

1) The consequences of making a mistake
2) **The lack of training data**
3) **The need for continuous adaptation**

1) The consequence of making a mistake

MA How many countries begin with the letter "Y"?

As of my last knowledge update in January 2022, there are three countries whose names start with the letter "Y." These countries are:

1. Yemen
2. Zambia
3. Zimbabwe

Please note that geopolitical changes can occur, and the number of countries in the world can change over time, so it's a good practice to verify this information with up-to-date sources for the most current and accurate count.

BBC

Home News Sport Business Innovation Culture Travel Earth Video Live

Man crushed to death by robot in South Korea

8 November 2023

Share

Emily Atkinson
BBC News

A man has been crushed to death by a robot in South Korea after it failed to differentiate him from the boxes of food it was handling, reports say.

Source: https://www.webfx.com/blog/marketing/chatgpt-fails/

Safety First

1X tests every humanoid in real-world scenarios before they're deployed. Their soft, organically-inspired mechanics make them safer from the inside out, so they're ready for your space.

Safe Intelligent Scalable

Source: https://www.1x.tech/

What makes learning robot skills harder than learning to master language (and vision)?

1) The consequences of making a mistake

2) **The lack of training data**

3) The need for continuous adaptation

The role of (labelled) data in Deep Learning

- **Deep Learning**: Breakthrough came from **Supervised Learning**
- E.g. collecting thousands/millions of images, with **human-made labels**
- Training neural networks to **generalize to new, unseen images**

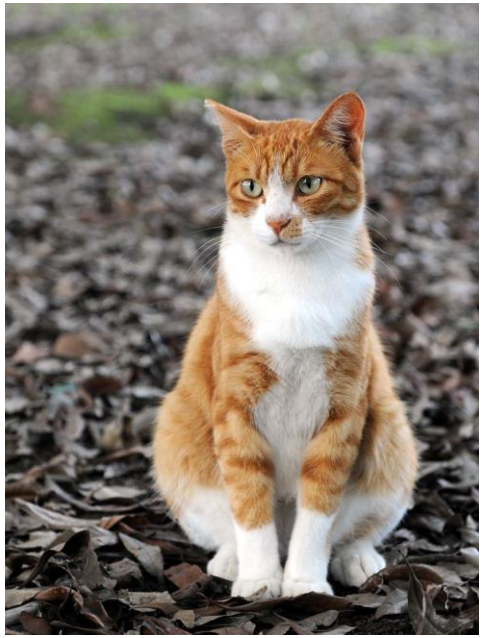

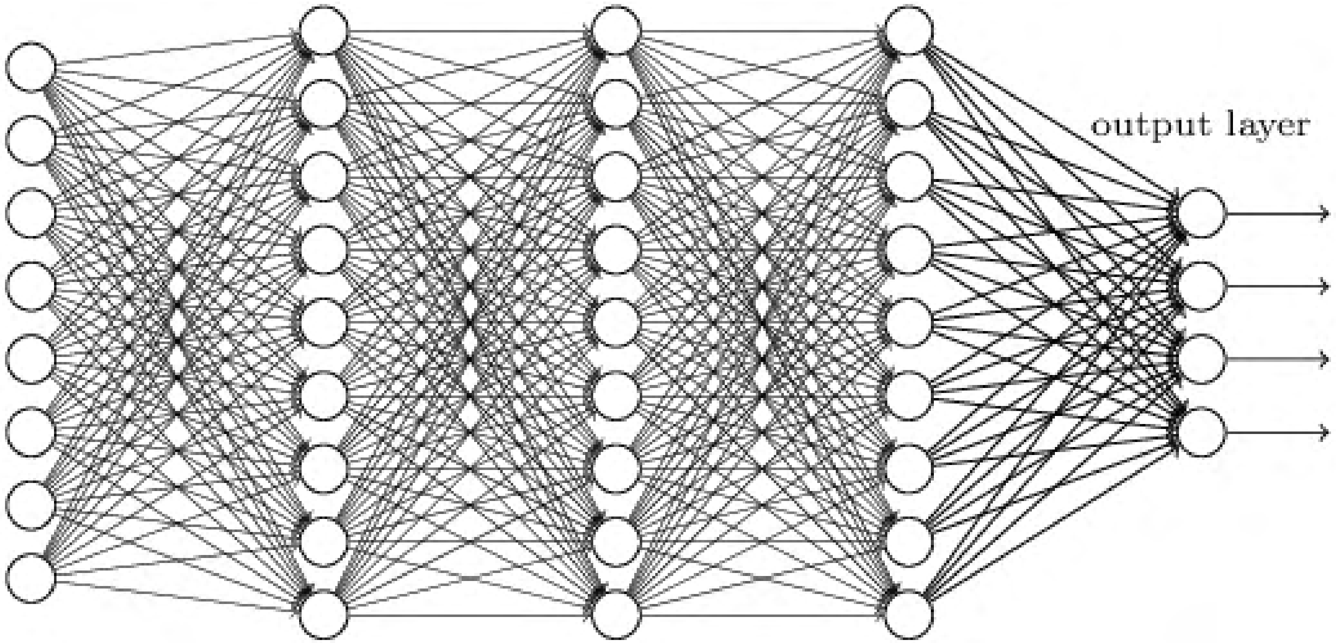

0.1	Katt
0.9	Hund
0.02	Esel
0.4	Elefant

But Large Models (e.g. GPT) require so much data we can't label it all!

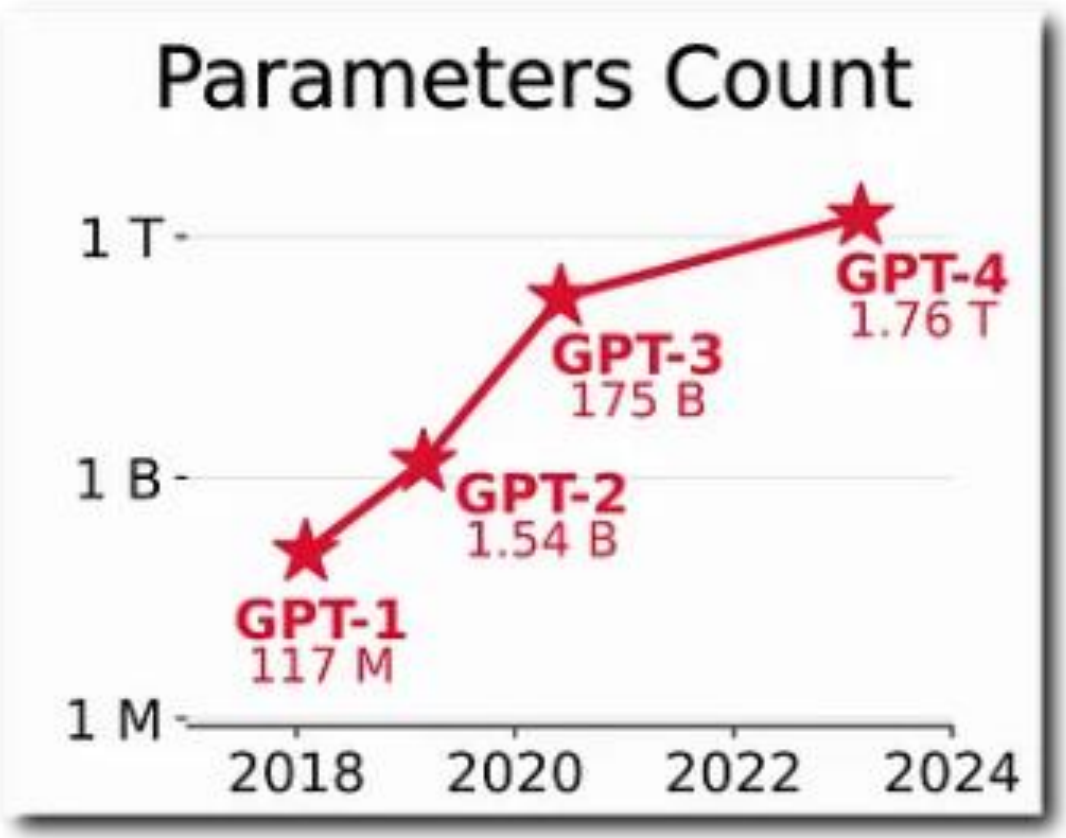

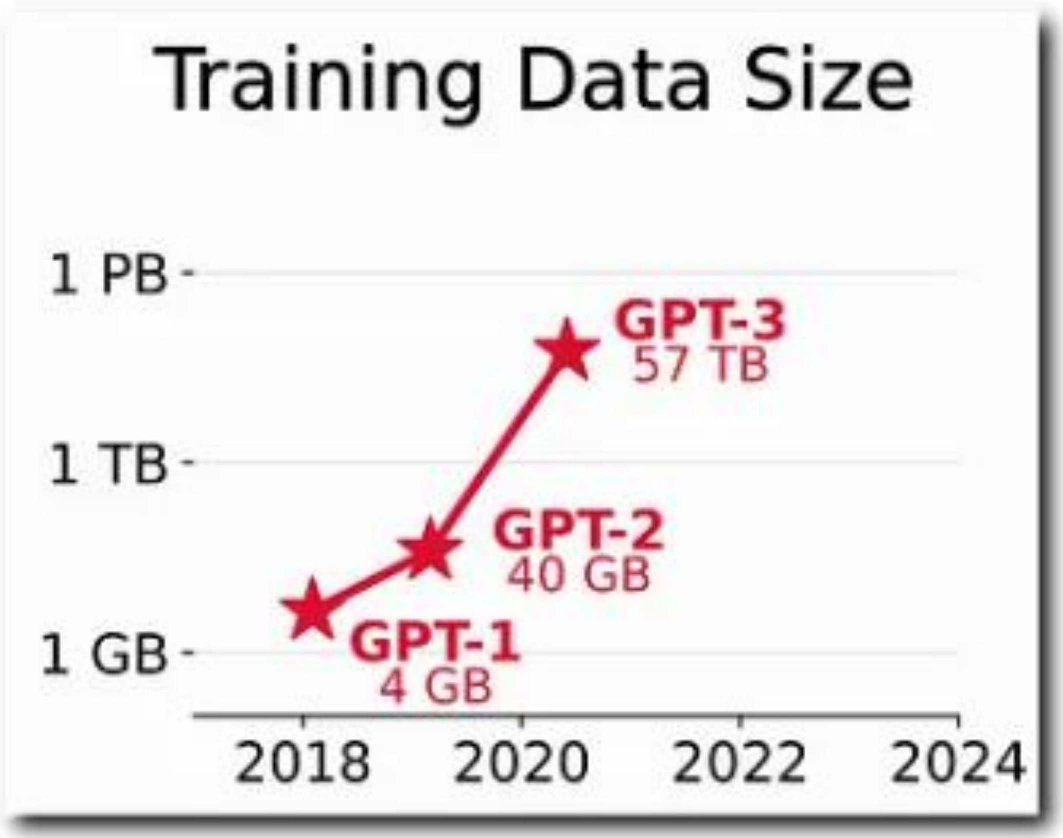

Source: https://medium.com/@francesco.casalegno/chatgpt-explained-whats-inside-the-ml-model-from-gpt-1-to-gpt-4-631d325368ef

Self-Supervised Learning

- **Large models** for language, images etc: Require so much data that **supervised learning doesn’t scale**
- **Solution: Self-supervised learning.** Get targets from the data themselves
  1) **Collect** huge amounts of **online data**, and
  2) Train a model taking **parts of the data as input,** and **another part as the target**

Self-Supervised Learning

- Text: Given a long text sequence as input – predict the most likely next word
- Images: Given a part of an image, predict what is in another part
- Video: Given 5 seconds of video, predict the next frame

Where would we get all this data for robotics?

A) **In a simulator**

- Reality gap

B) Human demonstrations ← Supervised robot learning?

C) Robot making its own demonstration data? ← Self-supervised robot learning?

Where to get data?
A) In a simulator

- **Model** your environment well in a **simulator**
- **Train** robot skills there (gives you lots of data without a big cost)
- **Transfer** learned skills to real world
- Deal with the **reality gap**

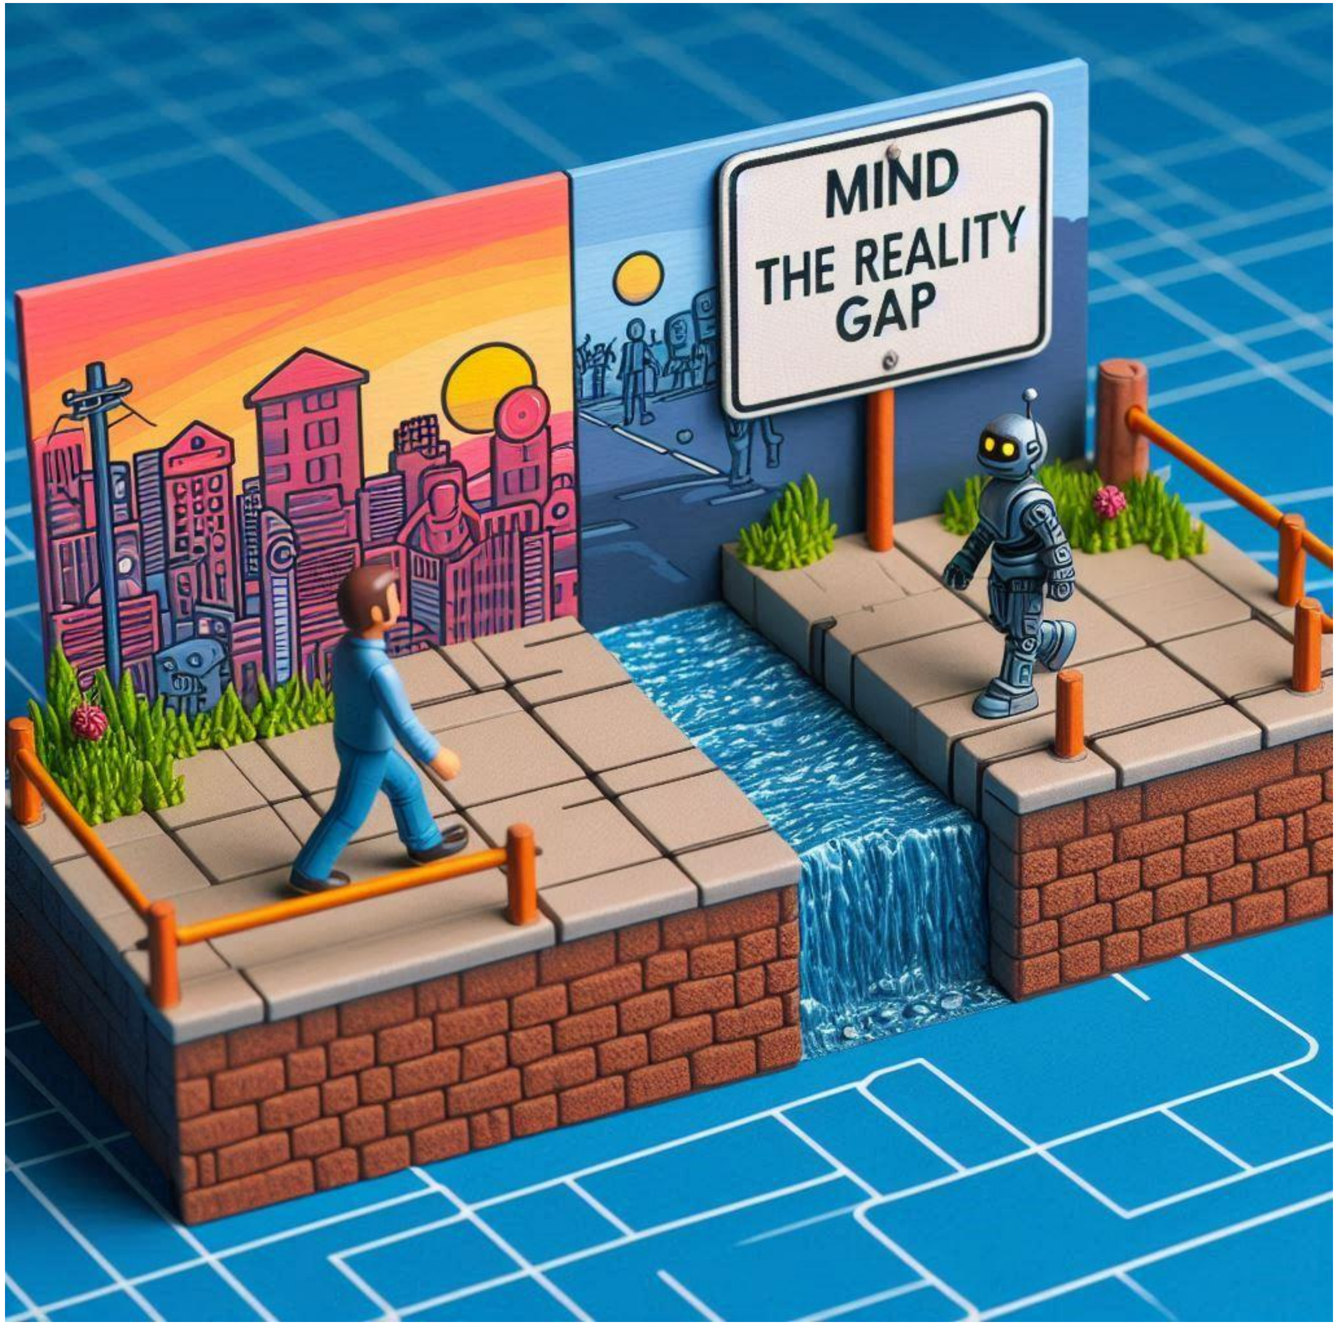

Quadratot: Reality gap

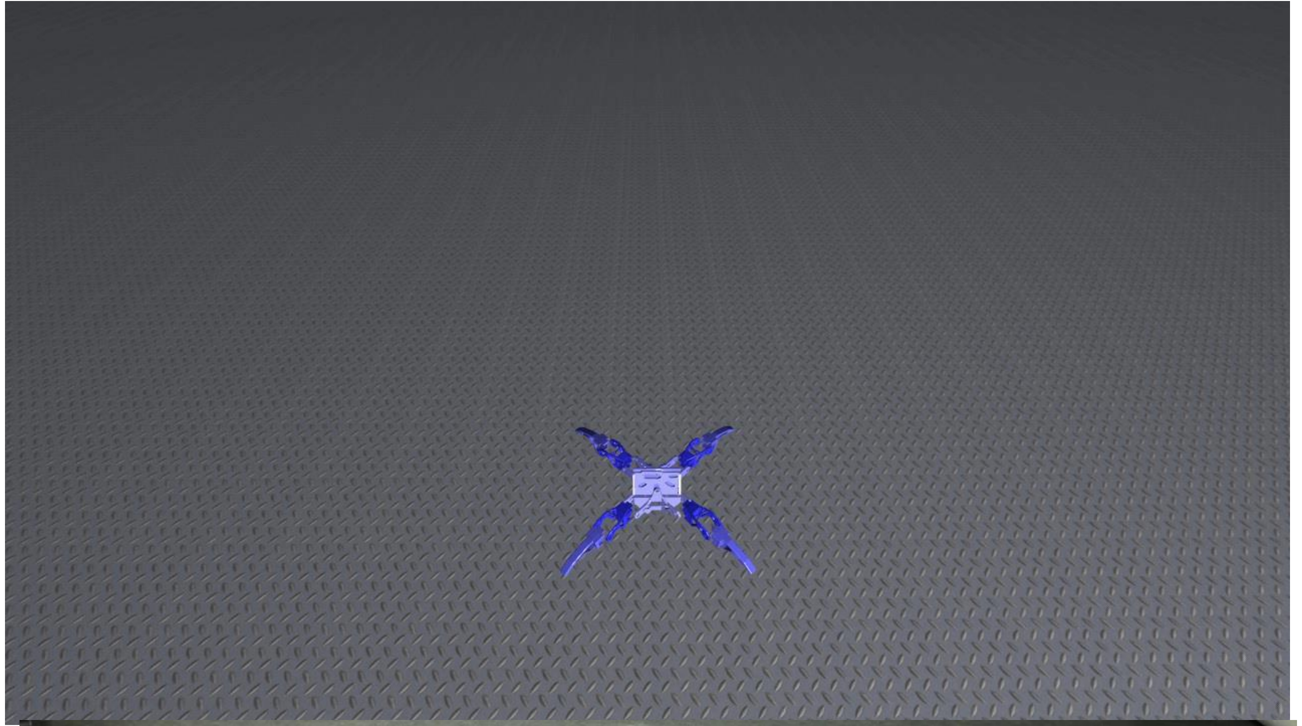

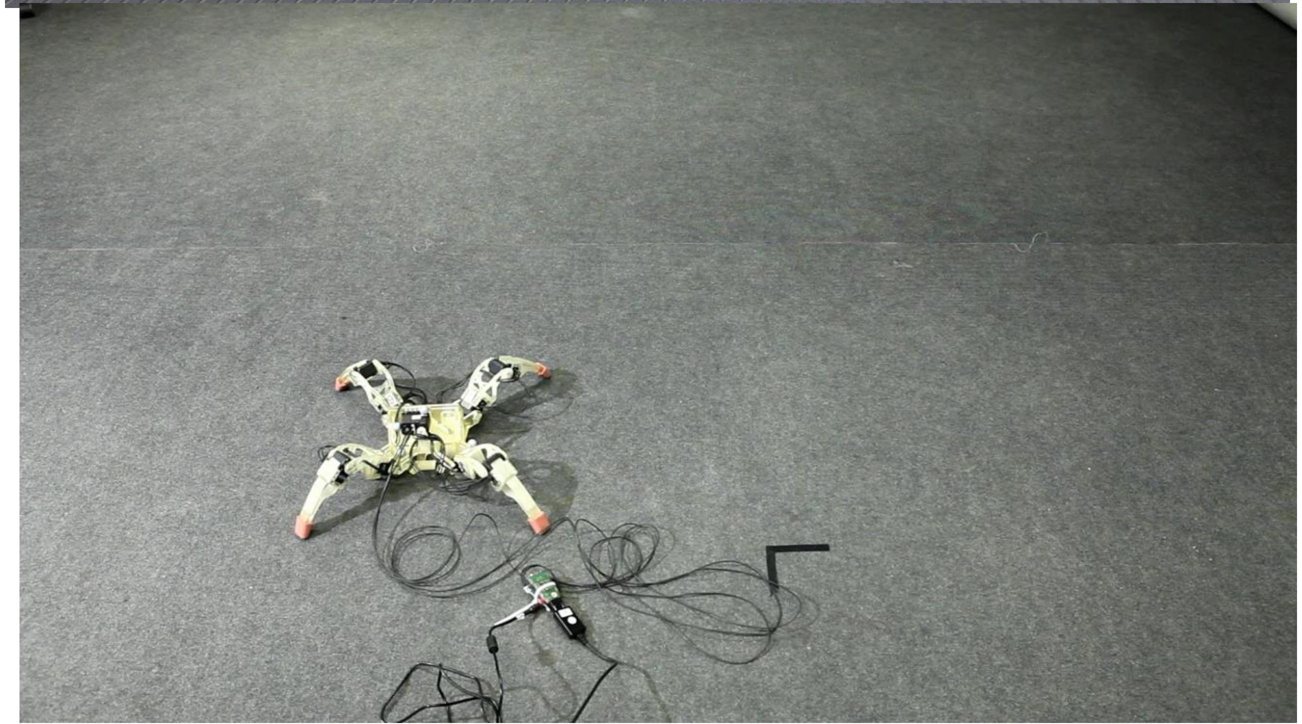

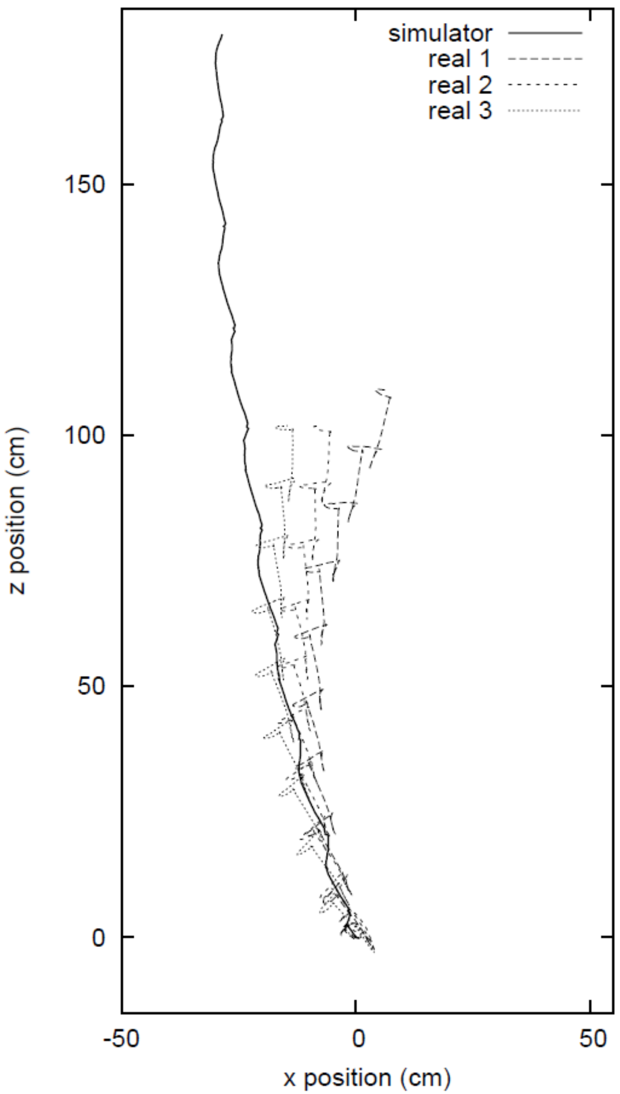

How can we reduce the reality gap?

- **Apply noise/randomness** during training
- Make **simulator more realistic** with **feedback** from real-world testing
- Try to **predict which solutions transfer best** to reality
- Let the robot **adapt/learn some more** after transferring to the physical world

Source: OpenAI

Where would we get all this data for robotics?

A) In a simulator

- Reality gap

B) **Human demonstrations** ⟵ Supervised robot learning?

C) Robot making its own demonstration data? ⟵ Self-supervised robot learning?

Where to get data?
B) Human demonstrations

SEARCH

FORTUNE

SIGN IN

Subscribe Now

Home News Tech Finance Leadership Well Recommends Fortune 500

TECH · TESLA

Tesla is hiring workers for $48 an hour to wear motion-capture suits to train its humanoid robots

BY SASHA ROGELBERG
August 19, 2024 at 7:29 PM GMT+2

Tesla announced the development of its humanoid robot Optimus in 2021.
STR/AFP—GETTY IMAGES

How many human demonstrations do we really need?

- One for each type of movement (pick, place, push, etc)?
- One for each type of object (toy, fork, apple, battery)?
- One for each combination of object and placement (toy on the floor, toy on a table, toy in a box)?

- The central question: How much can the robot **generalize** from the demonstrations it has seen to new, similar situations?

Where would we get all this data for robotics?

A) In a simulator

- Reality gap

B) Human demonstrations ← Supervised robot learning?

C) **Robot making its own demonstration data?** ← Self-supervised robot learning?

C) What if a robot could generate its own demonstration data?

- Could give similar advantages to **self-supervised learning, for robotics**
- **Planning based methods** can solve some robotics tasks
- But sometimes **planned actions fail** when carried out
- Can we **learn even from failed attempts** to follow a plan?

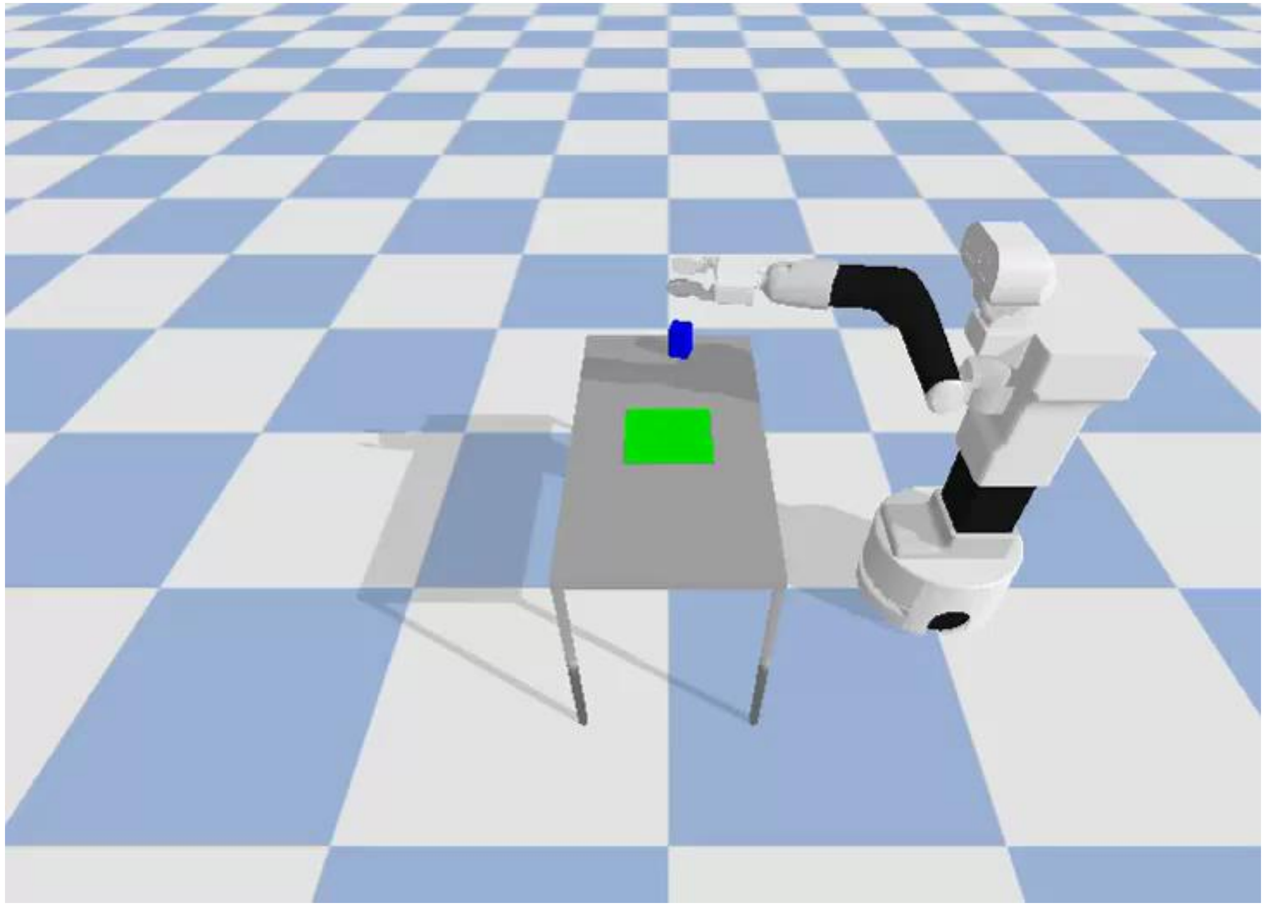

Improving Robot Skills by Integrating Task and Motion Planning with Learning from Demonstration

Shin Watanabe* Geir Horn Jim Tørresen Kai Olav Ellefsen

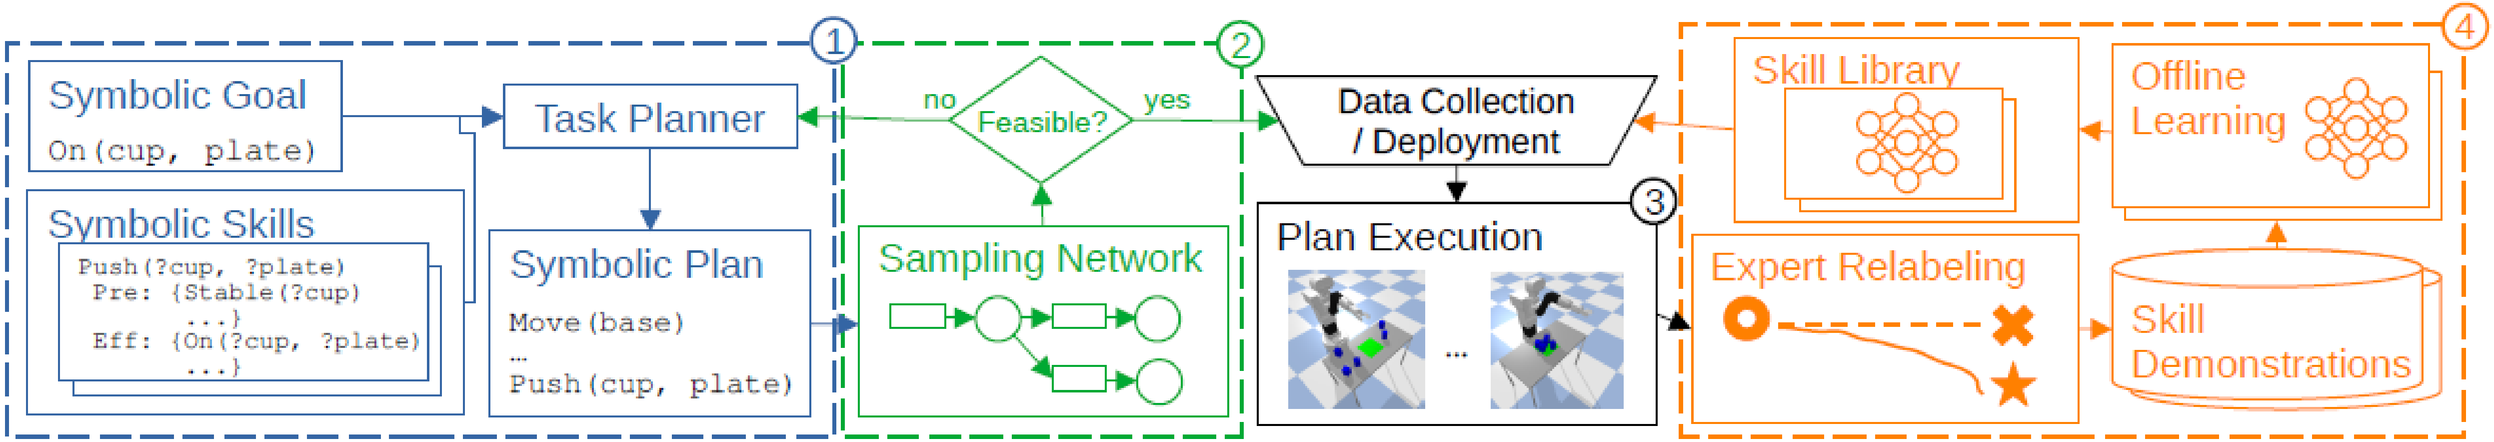

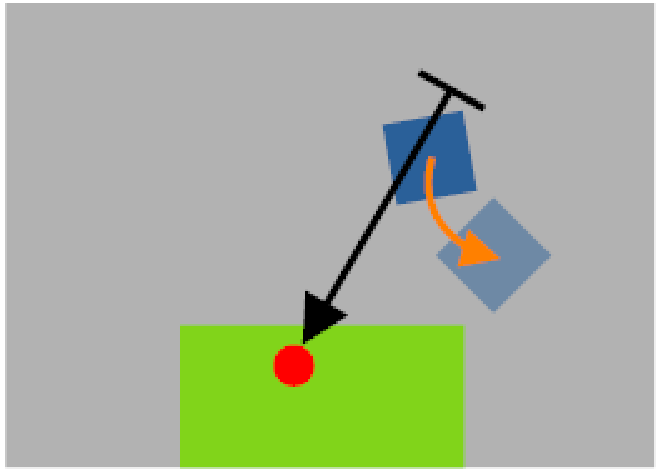

Limitation/future work

- This only works if we can learn something useful also when we fail.
  - Works for pushing, throwing, sliding objects
  - But not for grasping: If we fail to grasp, what can we learn from that?
- Can we find a way to do self-supervised robot learning for a wider range of tasks?

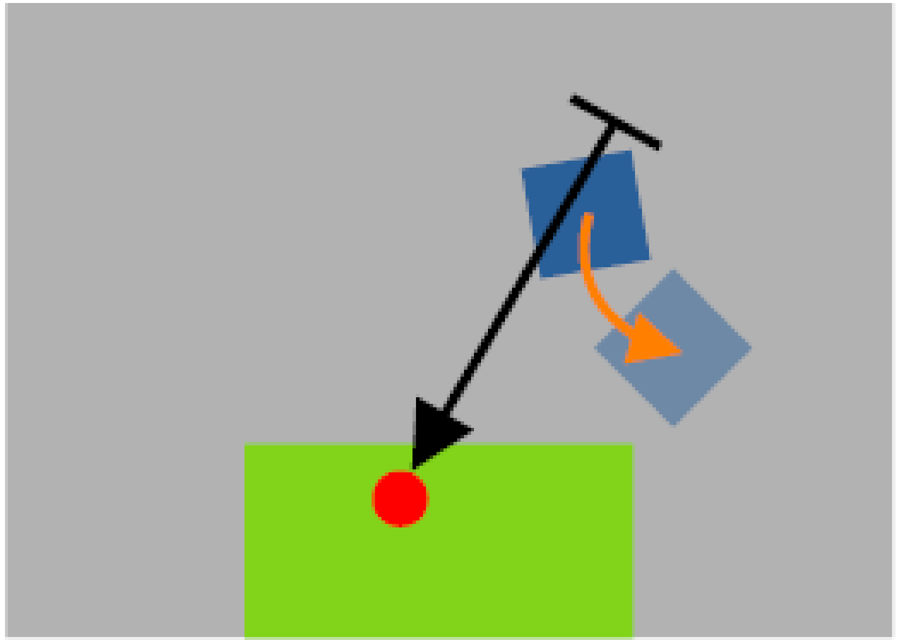

What makes learning robot skills harder than learning to master language (and vision)?

1) The consequences of making a mistake

2) The lack of training data

3) The need for continuous adaptation

3) Continual Learning

- There is a reason ChatGPT doesn’t learn directly from your inputs
- And that self-driving cars don’t update their software directly following an accident
- **Continual learning is difficult**
- But continuous learning and adaptation would be very useful for assistive robots

The way ML-models are typically trained

- ML-models are trained in **«batch» mode**: Collect lots of data, then train the model
- To train the **next version, collect more data and continue training** the model **with new and old data**
- 2 reasons this is not a good fit for robots

  1) We **don’t want to store all the data** from «old» experiences

  2) We would like to **be able to learn more flexibly** sometimes: Here is an important new experience. Integrate it with what you already know.

Continual learning in Artificial Neural Networks

- McCloskey and Cohen (1989):
  - –ANN learned ‘one’s addition facts
  - –Then learned ‘two’s addition facts
  - –Performance on ‘one’s addition decreased rapidly
    - Within 5 trials: 20 %
    - Within 15 trials: 0 %

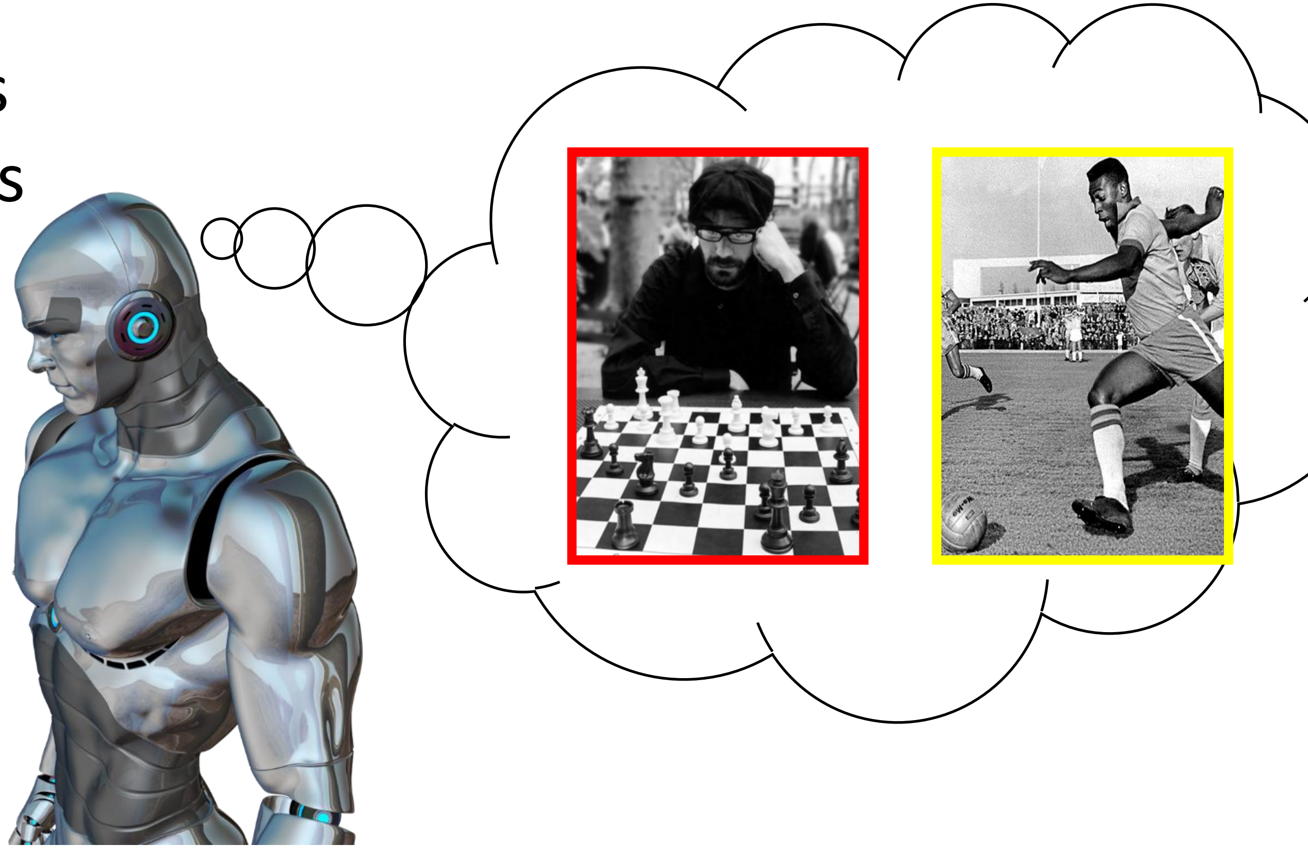

CATASTROPHIC INTERFERENCE IN CONNECTIONIST NETWORKS: THE SEQUENTIAL LEARNING PROBLEM

Michael McCloskey
Neal J. Cohen

Connectionist Models of Recognition Memory: Constraints Imposed by Learning and Forgetting Functions

Roger Ratcliff
Northwestern Universitv

Learning Skill A then Learning Skill B

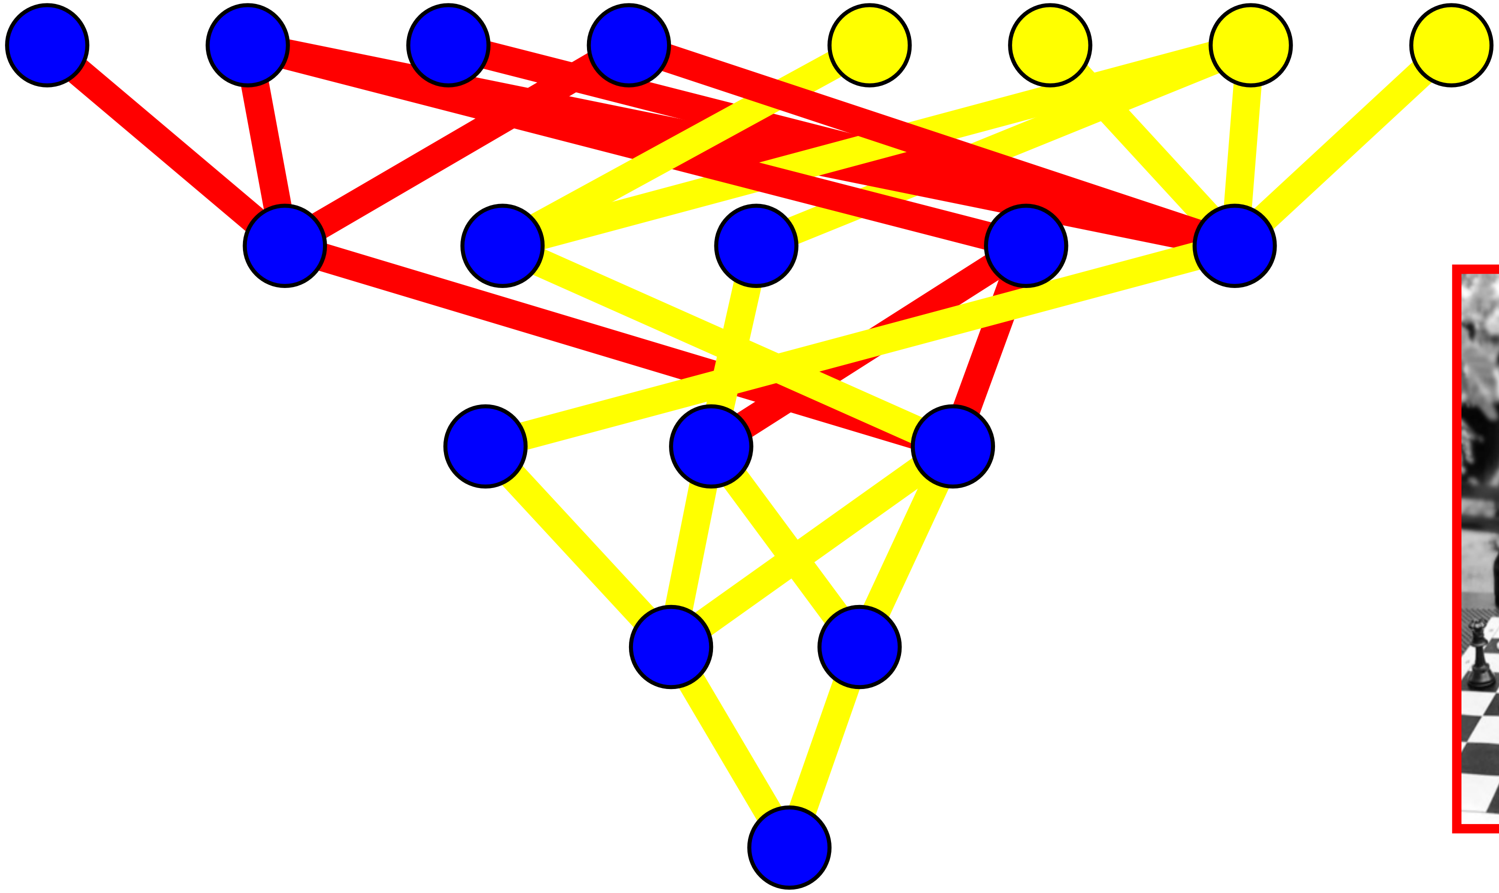

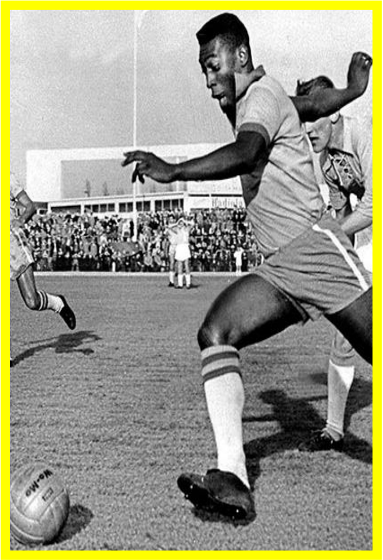

(Some of) The suggested Solutions

- **ANN structures that are less entangled**
- Rehearse/replay

PLOS COMPUTATIONAL BIOLOGY

OPEN ACCESS PEER-REVIEWED

RESEARCH ARTICLE

Neural Modularity Helps Organisms Evolve to Learn New Skills without Forgetting Old Skills

Kai Olav Ellefsen, Jean-Baptiste Mouret, Jeff Clune

Published: April 2, 2015 • https://doi.org/10.1371/journal.pcbi.1004128

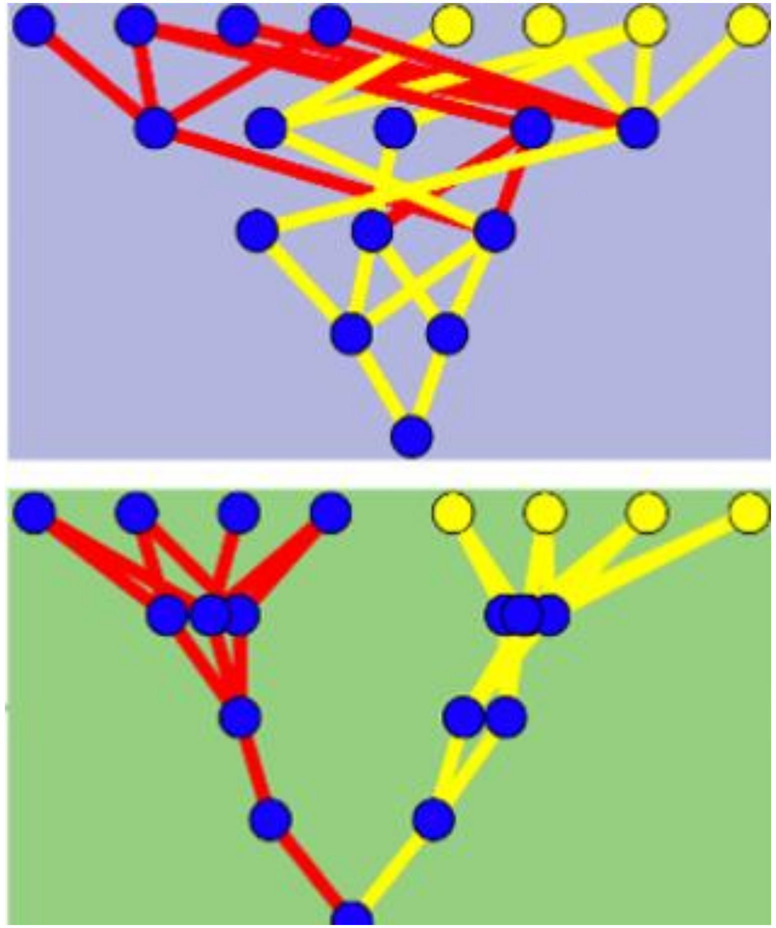

(Some of) The suggested Solutions

- ANN structures that are less entangled
- **Rehearse/replay**

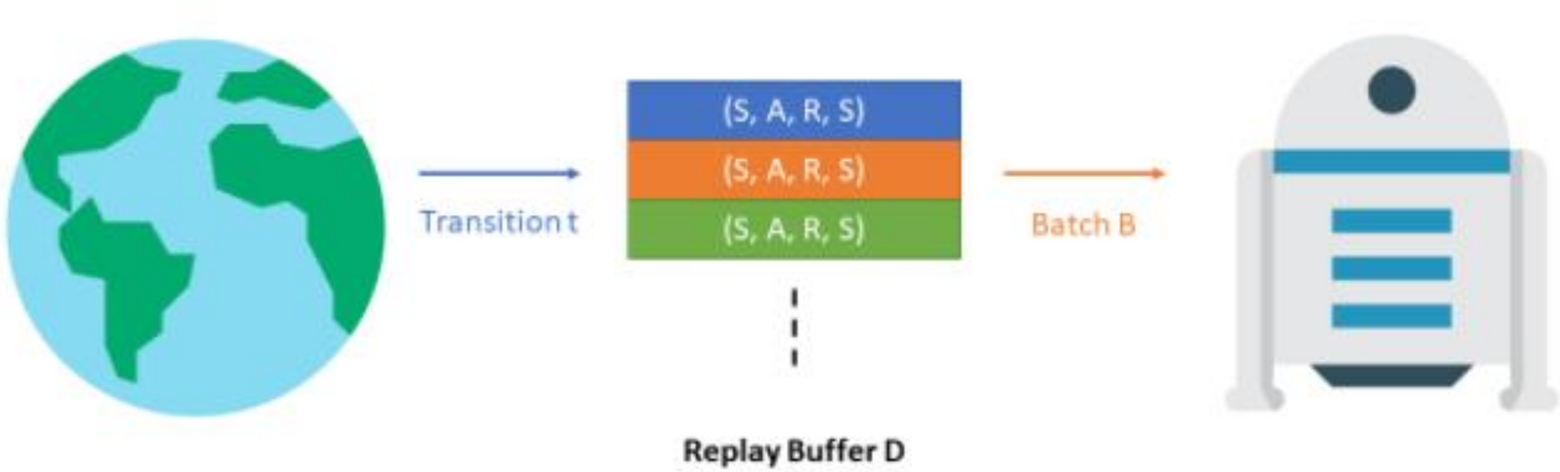

endtoend.ai

Exact replay

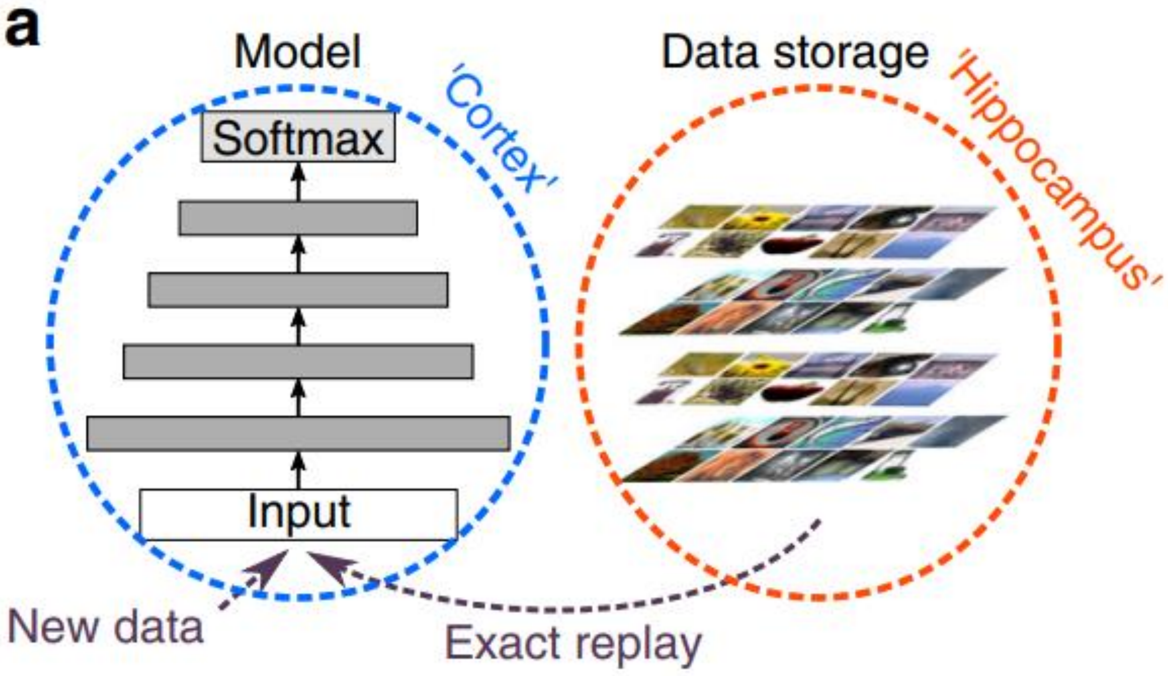


nature COMMUNICATIONS

ARTICLE

Check for updates

https://doi.org/10.1038/s41467-020-17866-2 OPEN

Brain-inspired replay for continual learning with artificial neural networks

Gido M. van de Ven[1,2], Hava T. Siegelmann[3] & Andreas S. Tolias[1,4]

What is the main limitation here?

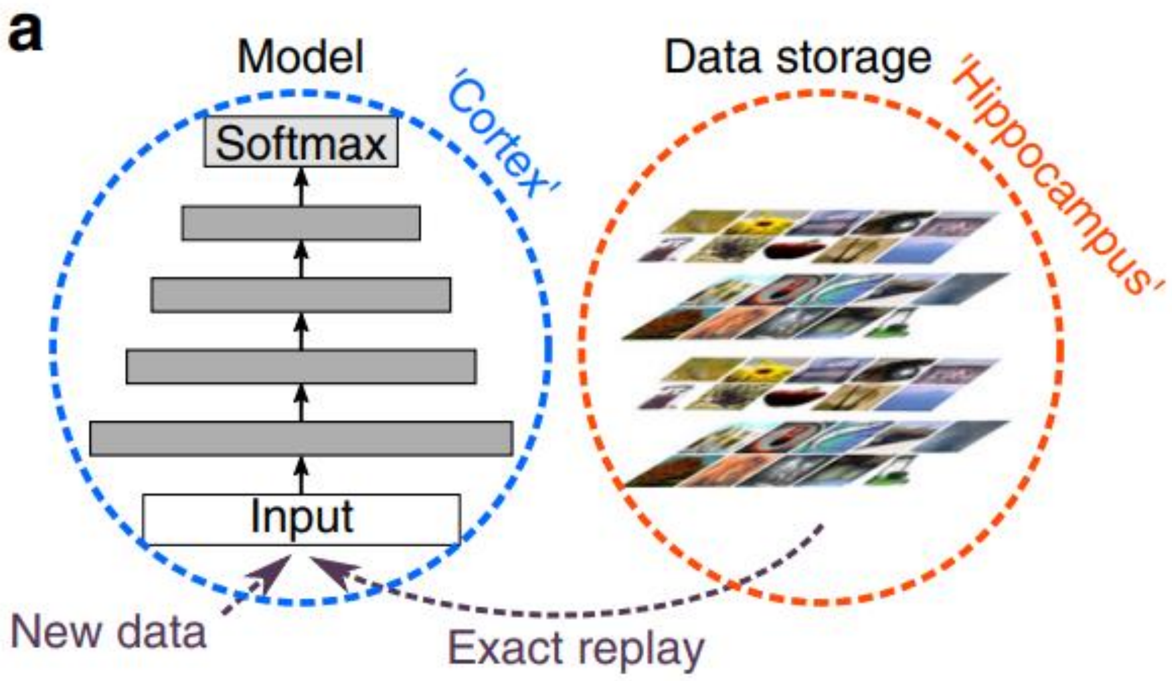

Generative Replay

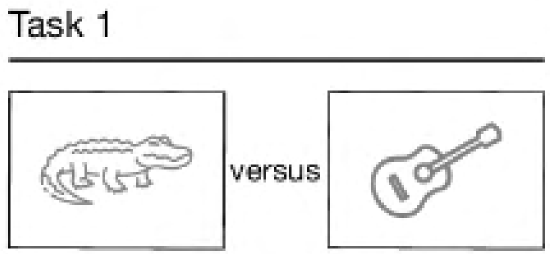

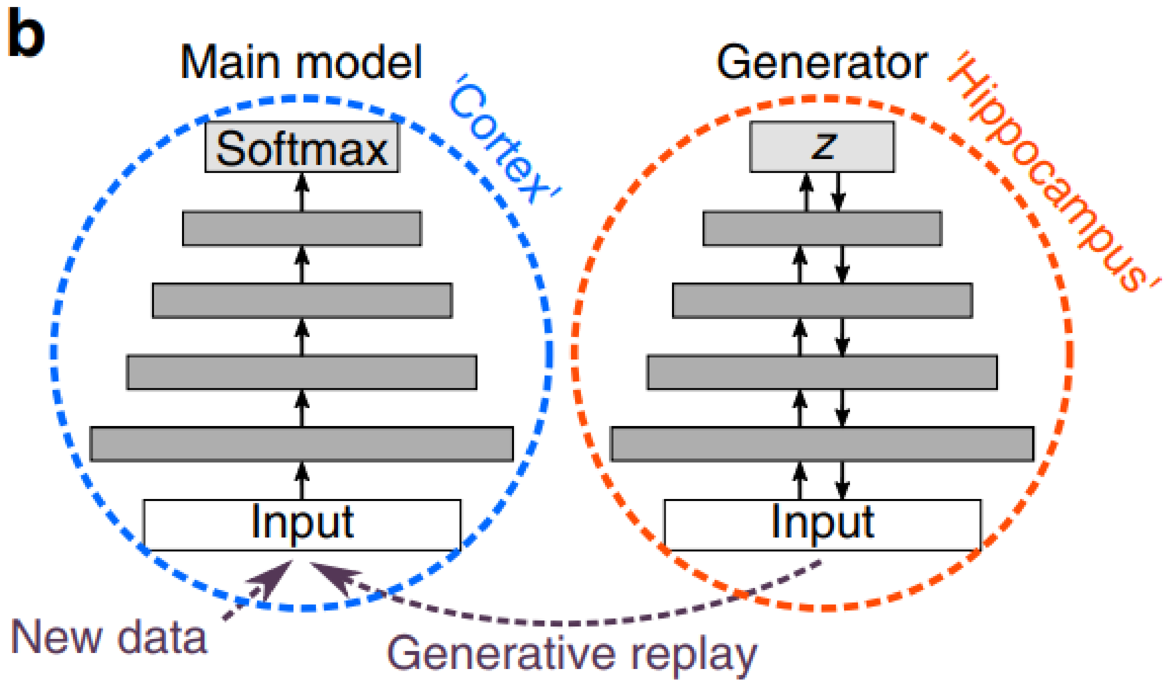

Generative Replay

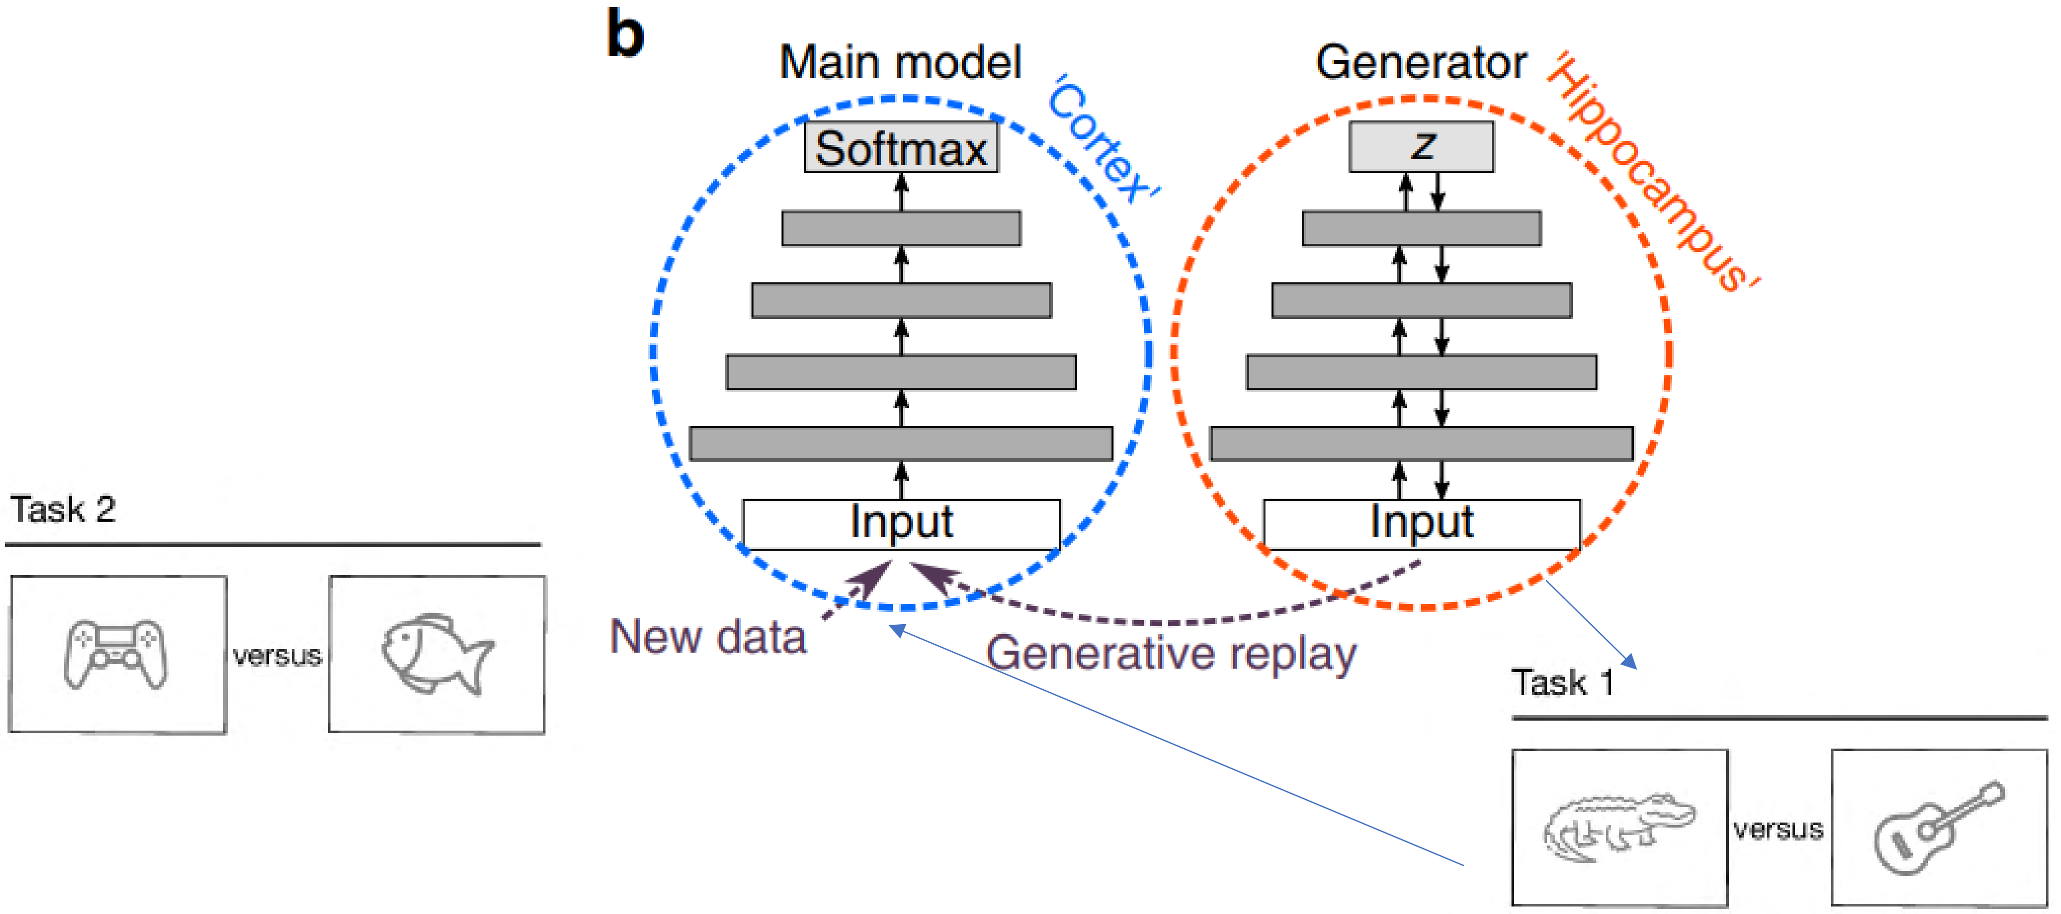

Generative Replay

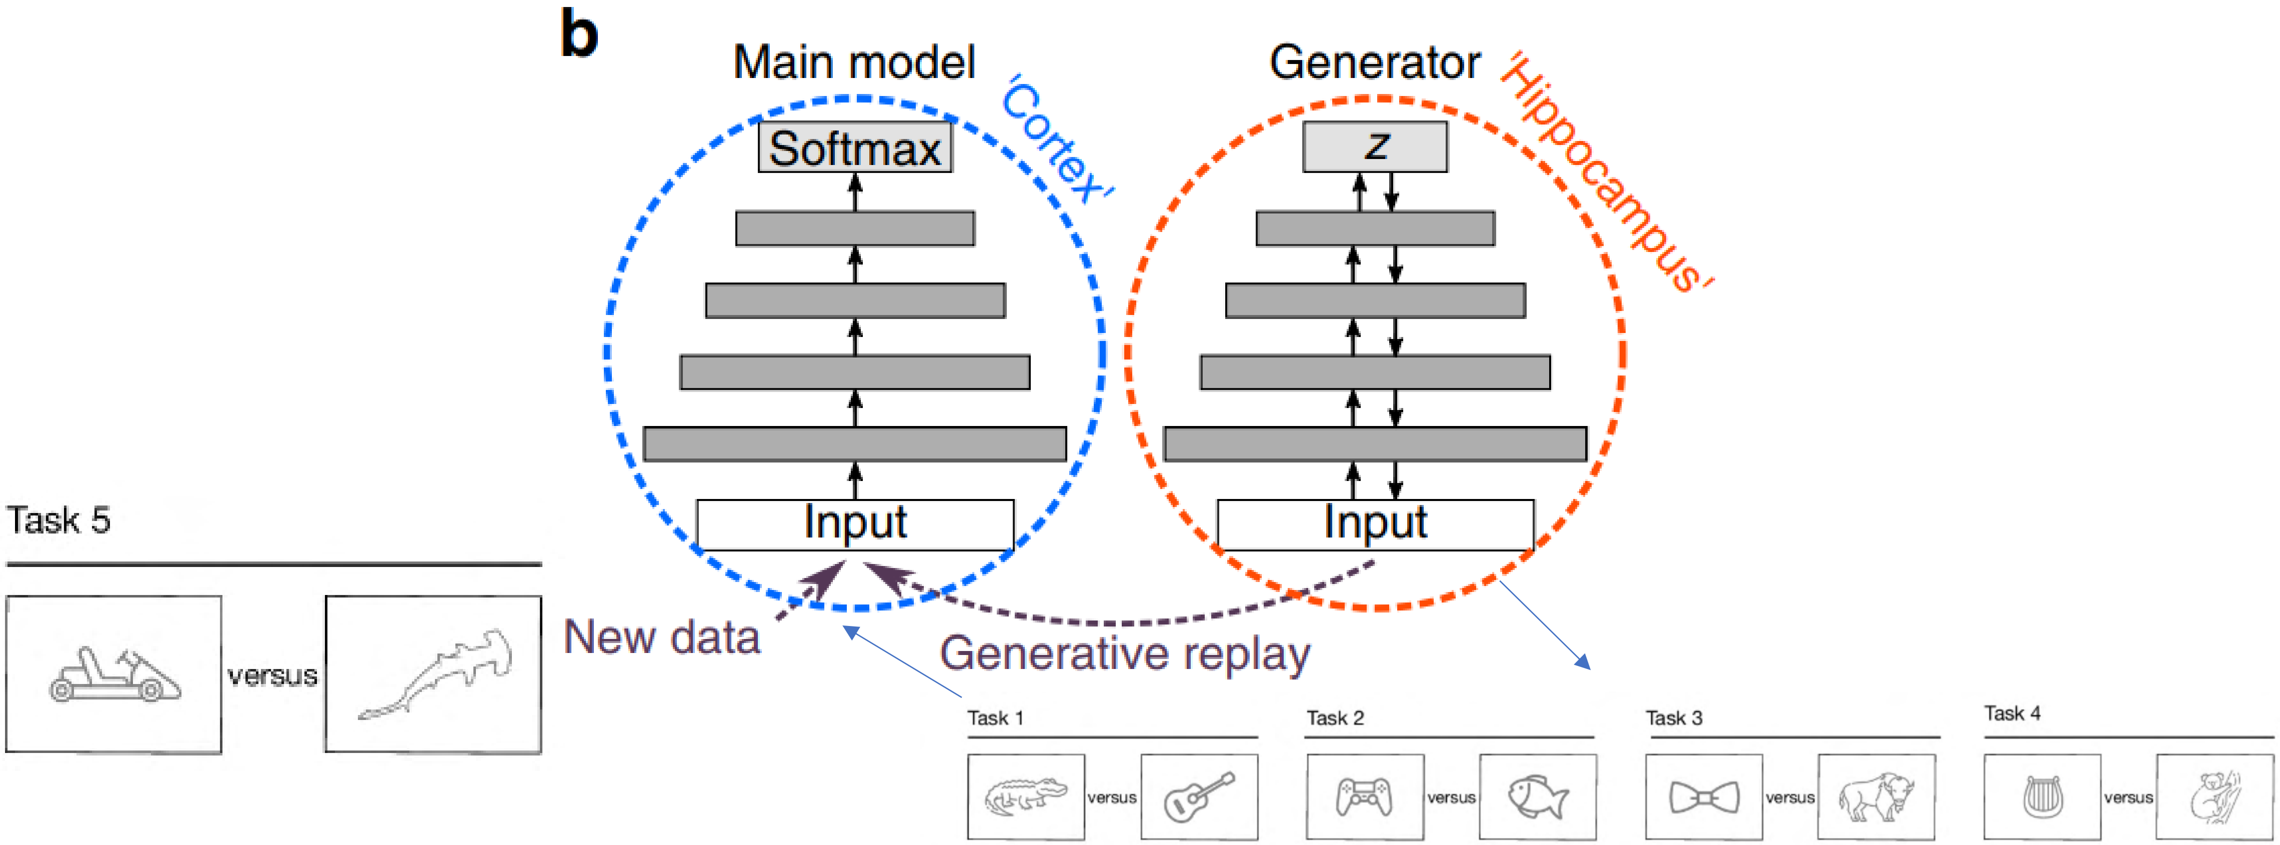

Problem: The generator can start to produce bad data

- The generator itself is trained continually
- This can lead to bad output, e.g. mixing up different classes
- That will in turn lead to bad training of the classifier
- **How to avoid this bad output?**

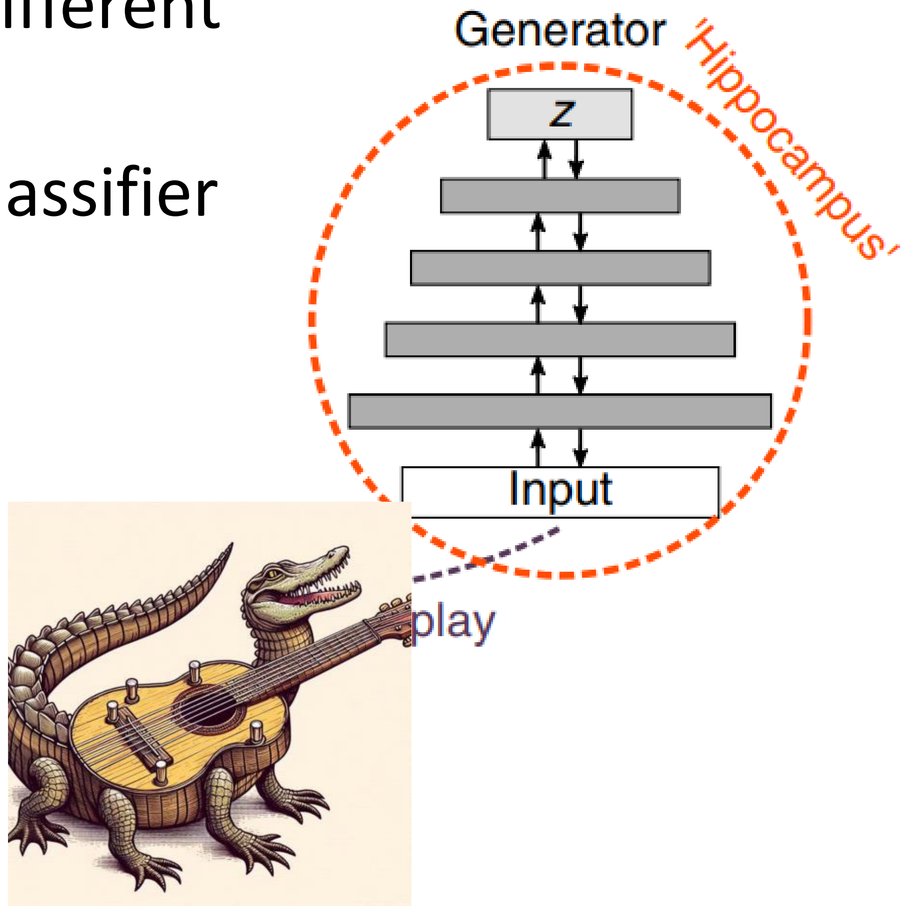

Possible solution: The classifier we are training can help us!

- We ask it: «How sure are you that this is belonging to a certain class»?
- If uncertain: We skip this, and generate some new data instead
- We call this **self-supervised generative replay**

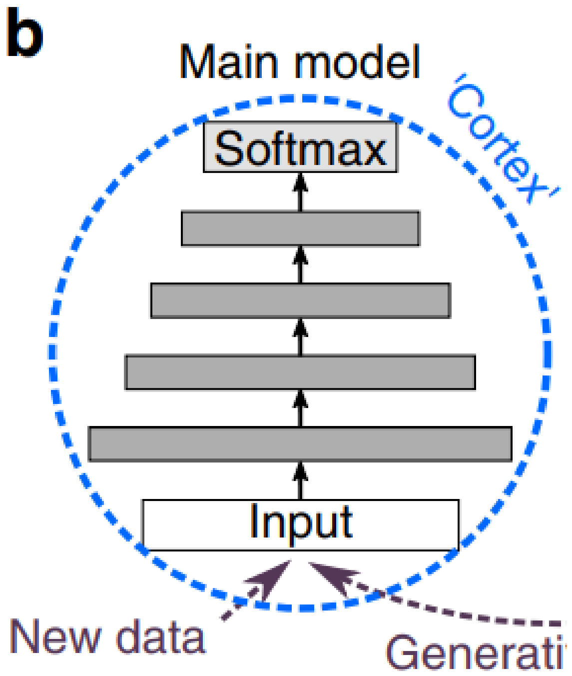

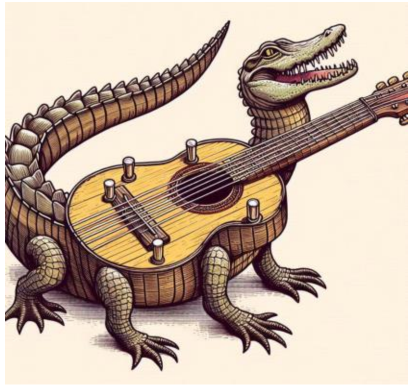

PhD Candidate Mateusz Wasiluk

Self-supervised generative replay

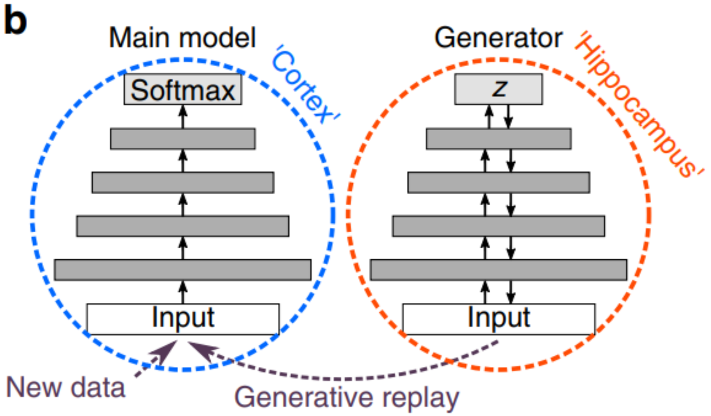

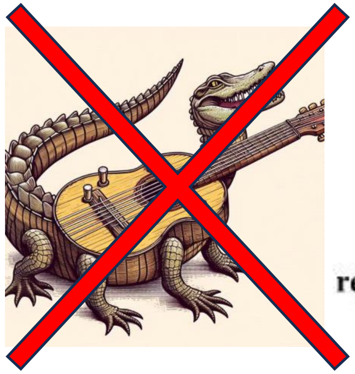

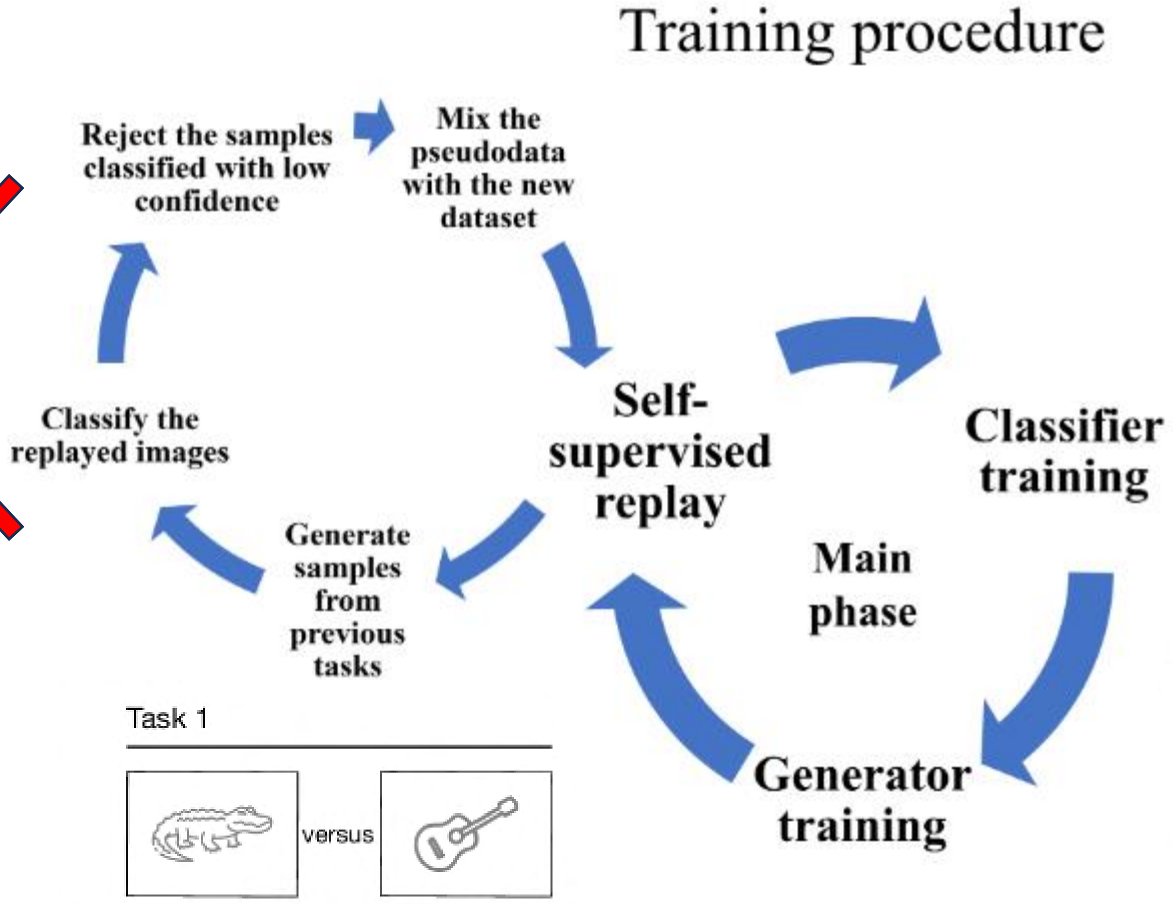

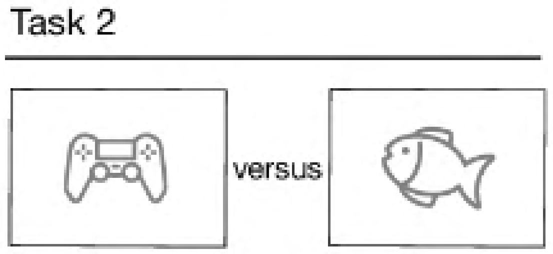

Preliminary result

- This filtering of bad generative data positively affects performance of the classifier
- The effect is larger the more tasks we add
- Challenge: Too strict filtering can reduce the diversity of the generator.

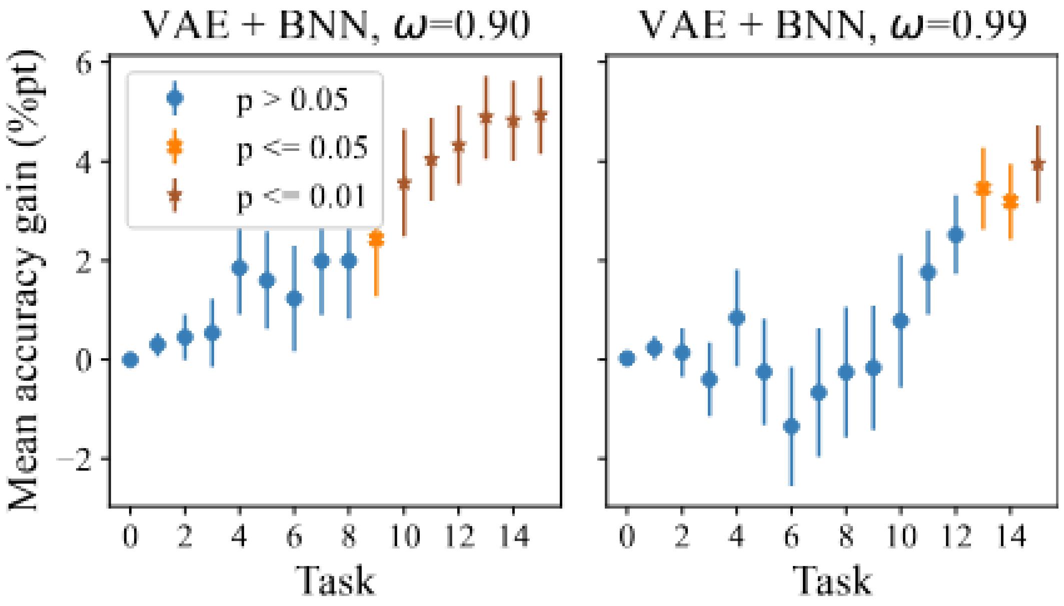

AI for robotics - conclusion

- **Big steps are being made**: Large datasets of demonstration data are being gathered, robotic «foundation models» are being trained
- We **don't know how far these will generalize** beyond their training data
- Some **fundamental challenges still need to be solved**, such as being able to learn new skills continously.